

PATENT APPLICATION
SYSTEM FOR PUBLISHING, ORGANIZING, ACCESSING AND
DISTRIBUTING INFORMATION IN A COMPUTER NETWORK

Inventor(s):

Alexandre Dayon
a citizen of France residing at
33 Rue Nicolo
75116 Paris
France

Assignee:

INSTRANET SARL
1 Square Chaptal
92309 Levallois-Perret
FRANCE

Entity: Small

SYSTEM FOR PUBLISHING, ORGANIZING, ACCESSING AND DISTRIBUTING INFORMATION IN A COMPUTER NETWORK

CLAIM OF PRIORITY

5 This application claims priority from Provisional Patent Application Serial
No. 60/148,029 filed on August 10, 1999, entitled "System for Publishing, Organizing,
Accessing and Distributing Information in a Computer Network," which is hereby
incorporated by reference as if set forth in full in this document.

COPYRIGHT NOTICE

10 A portion of the disclosure recited in the specification contains material
which is subject to copyright protection. Specifically, a Microfiche Appendix in
15 accordance with 37 CFR Section 1.96 is included that lists source code instructions for a
process by which the present invention is practiced in a computer system. The
Microfiche Appendix comprises 16 sheets of microfiche containing 1458 pages of source
code. The copyright owner has no objection to the facsimile reproduction of the
specification as filed in the Patent and Trademark Office. Otherwise all copyright rights
20 are reserved.

BACKGROUND OF THE INVENTION

25 The present invention relates in general to the transfer of data in computer
networks and more specifically to a system for publishing, organizing, accessing and
distributing information in a computer network.

30 Accessing information, and publishing information for others to access or
obtain, are important features of computer networks. However, although the trend is to
make information access and publishing easy for users of computers and computer
networks, many of the mechanisms available are not easy for an average user to master.
For example, publishing a web document not only requires a user to have some
knowledge about where to publish, and to what audience to publish, but the user may
have to publish the document to several "sites" or locations to make the document readily

available to a desired number of users. This is the case, for example, when a company uses the company's network, or intranet, that has different web sites associated with different departments, regions, etc.

5 The lack of structure or organization of web pages, and documents, on networks can be both good and bad. Lack of structure can allow easy publishing of documents without placing a burden on the publisher to comply with a predefined organization. This also lets each web site developer, online business, database, etc., to create a customized organization that is best suited to the specific type of information. However, lack of structure and organization also creates difficulties for a user of the
10 network to efficiently search for documents. Often a user has to perform many searches and access different websites and utilities to look for the document. This involves much typing and mouse (or other user input device) manipulation, is time-consuming and can be frustrating and counter-productive.

An analogy can be made to a newspaper which has an effective and well-
15 known organization. A reader of the newspaper can quickly obtain information from the newspaper by going to a subject section such as "Business," "Sports," "Travel," etc. The newspaper also provides an index or table of contents. Articles are organized in order of importance with "links" to other sections of related news, such as the continuing text of an article. However, to achieve this level of organization means that considerable time
20 must be spent on editing, page layout, paste-up, etc. Also, writers, editors, and other people must work in a concerted effort to produce the organized information. The approach of computer networks has been to allow each writer/publisher to throw an article into a haphazard network "bin" and to rely on loose organization mechanisms such as keyword searching, folder organization, hyperlink organization or criteria organization.

25 An example of a web site structure that a typical company might provide on its internal intranet is to have different web sites for functions, or departments, such as "Human Resources," "Marketing," and "Finance." If the company is large, there may be different regional offices, each having these functions. If, for example, the company has offices at locations in the U.S., Europe and Asia, this amounts to 9 different intranet sites
30 for information. Also, there will typically also be a main site for each regional office, a main site for each department or organization, and a main site for the overall company.

Thus, we find 16 possible sites in all for the example discussed above. The typical organization for documents associated with these sites is to have the documents pointed at by links. The links can be organized into categories at each site's

web page. Not only does this make publishing information extremely difficult when it is desired to make information available to more than one site; but any person interested in searching and obtaining information may have to visit several sites. Also, the task of publishing documents to the various intranet sites is usually handled by a different person than the writer/publisher. Not only can this become a huge task, given the number of documents and sites, but mistakes in classification are likely.

Some traditional methods for accessing documents in computer networks include keyword searching. This allows a user to make a relational query such as "movie review." The search will return documents that include the term "movie review" somewhere in the document. The documents can be at any number of sites. A search can be further narrowed by, for example, including a relational term such as "AND" and the name of a movie reviewer. Also, a specific date, or period of time, can be specified. However, because of the huge volume of information on most intranets (and certainly the worldwide Internet) the number of documents that match basic keyword searches is very large. Unless the user is very familiar with the terminology, and type, of documents relating to the subject in which the user is interested, the user's keyword search will most likely turn up many documents in which the user is not interested. These must be further filtered by refining the query until the proper documents are identified. With this approach it is often impossible to obtain a list of only relevant documents in which the searcher is interested. The scope of the keyword search can not be set by the user but is determined by the entity running the search engine and compiling the search engine database.

Besides the large volume of information, another difficulty in obtaining desired documents is that documents are created and "published" to the networks with few, or no, restrictions as to their form and organization. In other words, a web page can be created and published by a user that includes text, images, etc. with an arbitrary organization. A document might or might not have a title, author's name, publication date, etc. The text of a document can be arranged in columns, paragraphs, one-liner separated by images or graphics, etc. Often a document may not have any short identifying features, or any way to tell where one field, such as the subject of the document, begins and ends so that the subject may be indistinguishable from the body of the document at least insofar as a keyword search is concerned.

One approach to overcome some of these problems is to hand-annotate documents found on networks. Typically, this is done after document creation

(sometimes long after document creation) by a person who was involved in the creation of a document, web page, etc. Not only does this require substantial amounts of manual labor and time by persons having some skill and knowledge in the area to which the documents relates, but, by attempting to organize and summarize aspects of the document, mistakes can be introduced, thereby compromising the degree of accurate, searchable information.

Another approach is to use "folders" or sub-directories in programs such as email programs or web browsers. However, organizing information in this way is usually done manually by the viewer of the information (i.e., the email or web documents). There is no provision for publishing to a user's private organization as these folders are hidden from publishers. Where a public organizational hierarchy is implemented with folders, such a hierarchy often becomes large and complex, requiring much time to navigate. Also, this approach does not provide flexible security or access control.

Some web sites, such as www.yahoo.com, accumulate information such as documents, web pages, etc., from various sources and categorize, summarize and annotate the documents. This multicriteria organization defines categories which are presented to a user searching the Internet as a hierarchy of web pages. Each successive web page in the hierarchy (i.e., web pages progressively lower in the hierarchy) contain a new sub-category of selections that further narrow the category. At some point, the user decides that the category is the one desired and clicks a control. A collection of information that fits the category is then presented to the user.

However, this approach often requires that documents be interpreted and classified by a person other than the author so that errors can be introduced. Also, a considerable amount of work is required to do the classifying, write an abstract, , etc. Another drawback is that the navigation through web pages can be slow. Also, a user does not have an awareness of the overall classification scheme being used. In other words, the user does not know how many sub-category levels there are in the hierarchy, or what types of classifications are used, until the user has done a substantial amount of investigating into the hierarchy "tree" classes.

Still other drawbacks of the prior art include the inability to index to individual pages, sections, or portions of a document. This means that text that would otherwise be maintained as a single document must be broken into several documents if it is desired to only allow certain groups to have access to different portions of the original

text. Current network organizations do not provide a very flexible security and access system. Usually a website is restricted to user's with a certain account or password. Each user wishing to access the site, and all of the site's documents, must enter the password. The use of passwords is difficult to maintain since accounts must be set-up, user's can
5 forget the passwords, etc. Also, the granularity of password protection is very coarse as an entire website is usually either open or closed to a particular user.

Thus, it is desirable to provide a computer network-based system that overcomes some or all of the problems in the prior art and provides an efficient system for publishing, organizing, accessing and distributing information in a computer network.

SUMMARY OF THE INVENTION

The present invention allows a user to filter and view documents provided over a network. The documents are filtered based on categories. A category selector is a user interface tool that allows the user to select multiple values to define a category. For
15 example, values can be used to specify geographic location, corporate department, employee classification, time period, etc. Each category includes one or more values. For example, in the "Geographic Location" type of value, values can be "Worldwide," "Europe," "France," etc. Values for "Corporate Department" can include, e.g., "Human Resources," "Marketing," etc. Values for "Employee Classification" can include "All
20 Employees," "Mid-Level Managers," "Staff," etc.

Using the four value types discussed above, a category might be "Europe; Human Resources; All Employees; Before 1999." Thus, the category defines a segment, or filter, of documents or information within the total documents or information available in a network. The total documents or information available can be considered those
25 documents within a company's intranet, or the total of the documents available on the Internet, world wide web, individual web site, or other computer network or database.

Note that, as discussed below, two different users who have selected the same category may not see the same documents, or document sections. This is because the present invention takes into account the community to which each user is assigned.
30 So, for example, a user who has an HR dimension to their assigned community can be restricted from seeing documents that are designated for accounting. The same type of restrictions can apply to sections, or portions, of documents.

A preferred embodiment of the invention generates a list of documents that satisfy the category definition. The category definition is selectable by a user of the

system using a "selector" tool in a user interface. The selector tool in the preferred embodiment presents a list of values for each value type in the category. Thus, the user can choose, in a menu-like style, values to create a category definition. Once the category definition is set, only those documents which meet the definition are shown to the user.

- 5 This acts to greatly simplify the user's search through many sites, servers, libraries, etc. on the network. Keyword searching can also be used in conjunction with category selection to provide a powerful search.

10 The list of documents that meet the category can be organized as discrete documents, or can be organized as collections of document types. In other words, documents can be organized into document types such as "career," "demo," "legal," "policies," etc. Within each of these document types can be discrete document names, sub-types, etc. The collection of types, and the types' documents and sub-types, that are returned for a given category is referred to as a "theme." Categories are referred to as "communities," or "slices." Type values are referred to as "dimensions." Communities, dimensions, dimension values, and themes can be set by a system administrator. Users have great flexibility in using themes and have some ability to define new themes. A given theme can be associated with more than one slice. Users who publish documents to the network are associated with a slice, or community. The default for a published document is to be associated with a slice of the publishing user. However, the user can choose to associate a document with a different slice, or slices. Or, the system administrator can change the association of the document.

25 In a preferred embodiment, the invention provides a method for accessing web pages on a network, wherein the network is coupled to a server computer and a user computer operated by a user, the user computer including a user input device and a display device, the method comprising transferring a portion of a web page from the server computer to the display device over the network, wherein the portion of a web page includes a selector allowing the user to select one or more of the following categories: geographic location, corporate department, employee classification, time period; detecting a user's choice by receiving information generated in response to signals from the user input device to indicate the one or more categories chosen by the user; identifying one or more web pages associated with information that meets the chosen categories; and sending information about the identified web pages to the user computer.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 illustrates a selector tool of the present invention;

Fig. 2A illustrates strict filtering;

Fig. 2B illustrates ascending filtering;

5 Fig 2C illustrates descending filtering;

Fig. 2D illustrates the use of theme headings;

Fig. 2E illustrates a user interface mechanism for defining themes;

Fig. 2F shows a text search box used in conjunction with a slice selector tool; and

10 Fig. 3 illustrates a data model of the preferred embodiment.

DESCRIPTION OF THE SPECIFIC EMBODIMENTS

15 The present invention is to be embodied in a suite of software products created and distributed by Intranet, Inc., of Palo Alto, California. Various aspects of the invention are discussed below, followed by a description of hardware suitable for use with the invention.

20 Obtaining Documents

Fig. 1 shows screen display 100 of the user interface of the present invention.

In Fig. 1, selector 110 is a set of four pull-down menus or lists. Each of pull-down menu corresponds to one "dimension" used to define the "slice," or "community." Each of pull-down menus 102, 104, 106 and 108 correspond, respectively, to the following dimensions: geographic location, corporate department, employee classification, time period.

25 Screen 100 shows the display after the user has set values for each of the dimensions as: "U.S."; Human Resources; All Employees; All Periods." Thus, the slice, or community defined in selector 110 of screen display 100 is for documents associated with any employees, published at any time, in the companies U.S. offices relating to the company's Human Resources Department.

A list of the results of specifying the slice shown in Fig. 1 is to the left of the web-page at 120 of Fig. 1. This list shows two document type headings as "Career

and Policies.” Within the career heading, are three documents identified as “Employment Philosophy,” “Referral Program,” and “Training & Development.” Under the “Policies” heading, are “Employment Guidelines” and “Policies.” The list of documents at 120 is retrieved after the server automatically executes database search instructions according to the defined slice. The server begins executing the database search based on the defined slice when the user clicks on the “Find” icon at 122.

The user can view a selected document by clicking on the document name in the list at 120. The document is shown in the document viewing area at 130 of Fig. 1, below selector 110 and to the right of list 120.

In addition to document titles, uniform resource locators (URLs) can be shown in the list. By clicking on a URL, the user is taken to the corresponding web site, web page or other resource associated with the clicked URL.

Table I shows examples of possible values for dimensions.

In Table I, values for each of the dimensions used in selector 110 of Fig. 1 are as follows: the dimension “Geographic Location” can have values such as “Worldwide,” “Asia,” “Europe,” “France,” “U.K.,” “America,” and “U.S.” Naturally, any number of geographic locations can be specified and placed into the list. The preferred embodiment allows values for dimensions to be organized in a hierarchy. As shown in Table I, the value “Worldwide” includes values encompassed by it. These “child” values (to which the value “Worldwide” is the parent) are indicated by progressively indenting to the right. Thus, “Asia,” “Europe,” and “America” are each given one level of indenting to the right. Under the “Europe” value are associated child values of “France” and “U.K.” Under the value of “America” is the value “U.S.”

Geographic Location	Corporate Department	Employee Classification	Time Period
Worldwide	All Departments	All employees	All Periods
Asia	Human Resources	Staff	Today
Europe	Corporate	Engineers	This Week
France	Marketing	Mid Level Managers	Last Week
U.K.	Engineering	Senior Managers	This Month
America	Accounting	Executives	
U.S.		Board of Directors	

TABLE I

Values for the other dimensions are shown in Table I. Note that other selections and arrangements of dimensions, and dimension values, are possible. In the preferred embodiment, a system administrator is the one that sets up the dimensions, and dimension values, that can be used by users to define communities or slices. Thus, the system administrator is given control over the highest level of organization of documents in the network. Users use the dimensions, and dimension values to form slices to filter information and to obtain lists of documents.

Other arrangements for implementing selector 110 of Fig. 1 are possible. That is, each dimension need not be a pull-down list of items. For example, the complete lists of all values for dimensions can be constantly displayed. The use of pull-down menus however, makes selector 110 more compact. Thus, leaving more room for, e.g., displaying the text of document, or for other purposes. For example, an arbitrary text entry can be allowed so that a list is not even provided. This may be useful, especially for the "Time Period" dimension so as to allow a user to type in dates and times. The time period can also be set graphically with sliders, clocks and calendars, etc.

Another example of a slice using the dimensions and dimension value shown in Table 1 is the following: the slice "Europe"; "Marketing"; "Staff"; "This Week"; results in a listing of all documents generated in the company's offices in Europe concerning the staff in the Marketing Department for the present week that the document request based on the slice is being made.

Note that, once a slide is set, all of the user's subsequent browsing takes place within that slice. In other words, the user will not be shown, in the list at 120 of Fig. 1, documents or document headings, not associated with the selected slide. Themes which define which document headings, and documents, are shown in association with the slice are discussed in more detail below.

Thus, the user is able to "filter" documents in a way that makes sense within a corporate structure. In the preferred embodiment, filtering can be at different levels. The possible levels are "Strict", "" and "Descending". Other embodiments can include "ascending" filtering which is also discussed below.

Strict filtering means that only the documents that are associated with the given slice are identified to the user. Fig. 2A illustrates the results of a search using strict

filtering with the dimensions and possible dimension values shown in Table I. In Fig. 2A, only 4 documents in 2 categories meet the community “WorldWide; Marketing; All Employees; All Periods” with strict filtering.

In ascending filtering, all documents are listed that are associated with the selected slice, or which are associated with other slices that match the selected slice but that have one or more dimension values that are a parent-value of the selected slice’s dimension values. For example, using the geographic location dimension, the value “WorldWide” is a parent-value of the value “Europe”. Similarly, the value “Europe” is a parent-value of the value “France”. Thus, if the user is filtering the slice shown in Fig. 1 of “U.S.; Human Resources; All Employees; All Periods” all documents associated with the following slices will also be returned: “America; Human Resources; All Employees; All Periods”, and “World-wide; Human Resources; All Employees; All Periods”.

A preferred embodiment of the invention uses only strict and descending filtering. Other embodiments can allow the user to select the type of filtering from among one or more of the three types. Other types of filtering are possible.

Fig. 2B shows an example of the list of documents returned with ascending filtering used on the community “Europe; Sales; All Employees; All Periods”. Documents are returned that are associated with both the European and worldwide sales reports.

“Descending” filtering is similar to “Ascending” filtering, but proceeds in the opposite direction with respect to dimension values. In “Descending” filtering, all documents associated with the selected slice, and slices having the same dimension values as the selected slice and including slices having a dimension value that is the child-value of one or more of the selected slices dimension values, are provided to the user.

For example, in the case where the slice specified is the “America; Human Resources; All Employees; All Periods” then documents associated with that slice and also the slice “U.S.; Human Resources; All Employees; All Periods” are also returned to the user.

Fig. 2C illustrates descending filtering. Fig. 2C shows the results returned with the community “WorldWide; Human Resources; All Employees; All Periods”. Note that many documents are identified which correspond to all geographic regions for Human Resources departments.

A refinement to the selector tool is shown in Fig. 2F.

In Fig. 2F, selector 140 includes dimensions similar to selector 110 of Fig. 1, but also includes the added dimension of “Hobbies.” Further, text entry window 142 is

provided, along with search button 144. The use of the text entry window and search button allows a user to perform a keyword search within the defined slice. All documents within the slice filtering which also include the text phase or keyword specified in the text entry window will then be displayed to the user. Note that refinements are possible, such as allowing relational expressions within the text entry window.

Thus, slices are used to provide a filtering function across many documents in one or more web sites, servers, databases, etc. The slice is easily chosen by the user with the selector tool described above. Different filtering modes can be employed.

An additional feature of the invention allows users to predefine slices that can be recalled later. Such pre-defined slices are referred to as "channels." Thus, a user can select a slice such as "U.S.; Human Resources; All Employees; All Periods" and select a label to associate with the slice, such as "All US HR." Once designated as a channel, a hyperlink named "All US HR" appears on the user's display. This can be on a web page or part of a persistent toolbar, sidebar, etc. A user can then conveniently invoke the slice by merely clicking on the channel. The user can define multiple channels in this manner.

Predefined channels, or lists of channels, can be prepared and sent to other users. For example, where the channels are listed on web page, the web page can be emailed to one or more other users. By opening the emailed web page and clicking on the desired channel, the slice associated with that channel becomes the user's selected slice.

Publishing

Users can publish documents to the system. The publishing, and accessing, of documents is best understood using the two concepts of "coordinate" location and "scope" of visibility.

dimension value ← A coordinate is the specific set of dimension values associated with a user or document (or document section). Another way to think of a coordinate is as the community to which a user or document is associated. For example, a user can belong to the "France"; "Human Resources" community. Typically, a document published by a user is given the same coordinate, or community, as the publishing user. A user with a coordinate of "France"; "Human Resources" would publish to that coordinate. However, the system can be set up by the system administrator so that other coordinates are used as

the default coordinates for publishing for any given user, department, geographic region, community, etc.

5 The user can also be permitted to publish the document to multiple coordinates. One way to do this is to allow a user to override the default settings by using selector 110 of Fig. 1 to create an association between the user's published document and any desired coordinate. The user publishes to a coordinate by having the document displayed in document viewing area 130 and by clicking on the "Publish" icon 124. At the time of clicking on Publish icon 124, whatever coordinate (i.e., set of dimension values) selected by selector 110 is associated with the document. A user publishes to multiple coordinates by repeatedly selecting the next coordinate to publish to and then clicking on the "Publish" icon at 124 while the document to be published is displayed. Note that the search rules described above, for example ascending filtering, permit users at different coordinates than a documents published coordinate to view and access the document as long as the users are within the filtering rule.

15 The notion of "scope" refers to the collection of all user coordinates that can view and access a document. This can be more than the users at the document's published coordinate and at coordinates included within the filtering rules.

20 The system administrator can add coordinates to a document's (or document section's) scope. The system administrator can authorize, and exclude, users from accessing predetermined coordinates. This provides security and access controls based on any of the dimension values such as geographic region, employee position, etc. The system administrator, or another user, can also exclude users belonging to certain coordinates from viewing certain documents, or document portions. Thus, a document, or document portion , as explained below, can make use of its assigned coordinates, within the system to establish the document's "scope of visibility" to users. Naturally, any schemes for permitting or restricting documents, or document portions, from viewing can be used. For example, an access list of each specific community, or coordinate, that can view a document can be maintained. An exclusion list can similarly be maintained in tandem with the access list. Portions of a document can each be associated with one or more coordinates which are different from coordinates associated with other portions of the same document.

When a document is created, it is considered as a single section. This is referred to as the "Master Section." In many situations, this is the only section of the document and acts to associate the entire document with a coordinate or theme (discussed

below). As the document is created, the document author can define document sections which can, in turn, be associated with different coordinates, themes and access requirements. For example, a document may include sales reports which are excluded from all communities that do not include the "sales" or "executive" values.

5 Sections of a document can include, or entirely be, executable code such as Java, XML. Sections of a document can provide other mechanisms to invoke streaming media information, or perform other tasks. Such executable sections can be associated and restricted in the same way as text sections, described above. This allows greater flexibility, capability and control. For example, an executable section can cause a video or ShockWave™ window to appear and begin playing. Java can be used to have an
10 embedded calculation box, spreadsheet, etc. appear in a document. Such functionality can be excluded from users not associated with predetermined slices, or from users who do not wish to access the functionality. Other examples include a price list where the European prices are only visible by people in European communities and where the U.S.
15 prices are only visible by people in the U.S., etc.

Themes

 One aspect of the invention uses "themes" to help organize the presentation of information to a user, and to assist in publication of information across
20 communities. Themes function as category headings for groups of documents which fall
 into the category.

 Fig. 2D shows a portion of the user interface of the present invention to illustrate theme headings. In Fig. 2D, theme headings are shown at 150. Below each heading are document titles. The document titles are indented from their respective
25 theme headings. Although not shown in the black-and-white Figure, theme headings are red in the preferred embodiment. Theme headings function as standard category headings. For example, under the "Demo" theme heading, the documents "Our Products," "Kit PreSales" and "Kit Sales" can be found. Although not shown, sub-themes can be used. For example, the "Demo" theme can have a sub-theme of "Kits"
30 under which the last two documents would be listed.

 Unlike communities which serve to categorize, and distinguish, documents according to company structure and timeframe; themes organize documents by document subject or type. In other words, dimensions describe a group of people, or community, while themes define content classification. For example, a theme can be documents

category headings for documents
document types
subject content classification

related to employee "Stock Options." Such a theme can be associated only with countries that allow ownership of stock. In the example dimensions and values used so far, such a theme would be associated with the "U.S." and "U.K." values of the "Geographic Region" dimension. A publisher can publish to this theme by selecting from a list of themes at the time of publishing. The published document is automatically associated with the U.S. and U.K. values. The theme heading will only show up in communities that include the U.S. and U.K. values (assuming strict filtering), along with any documents associated with the theme. Users can use themes as part of a filter so that only documents matching selected themes' subject categories (in addition to the dimension values of the slice, or community) will be displayed.

In the preferred embodiment, a theme is selected for a document by using a pull-down list such as that shown in Fig. 2E. The user associates a theme with a section by placing the cursor within the section (e.g., by clicking in a text region of the section) and by selecting the theme. Each user is assigned a default community to which the user belongs. A theme for the document is assigned by placing the cursor in the Master Section and selecting on the desired theme.

Hardware/Software Implementation

Figs. 2A-C illustrate basic hardware suitable for use with the present invention.

Fig. 2A shows computer system 1 including display 3 having display screen 5. Cabinet 7 houses standard computer components (not shown) such as a disk drive, CDROM drive, display adapter, network card, random access memory (RAM), central processing unit (CPU), and other components, subsystems and devices. User input devices such as mouse 11 having buttons 13, and keyboard 9 are shown. Other user input devices such as a trackball, touch-screen, digitizing tablet, etc. can be used. In general, the computer system is illustrative of but one type of computer system, such as a desktop computer, suitable for use with the present invention. Computers can be configured with many different hardware components and can be made in many dimensions and styles (e.g., laptop, palmtop, pentop, server, workstation, mainframe). Any hardware platform suitable for performing the processing described herein is suitable for use with the present invention.

Fig. 2B illustrates subsystems that might typically be found in a computer such as computer 100.

In Fig. 2B, subsystems within box 20 are directly interfaced to internal bus 22. Such subsystems typically are contained within the computer system such as within cabinet 7 of Fig. 2A. Subsystems include input/output (I/O) controller 24, System Random Access Memory (RAM) 26, Central Processing Unit (CPU) 28, Display Adapter 30, Serial Port 40, Fixed Disk 42 and Network Interface Adapter 44. The use of bus 22 allows each of the subsystems to transfer data among the subsystems and, most importantly, with the CPU. External devices can communicate with the CPU or other subsystems via bus 22 by interfacing with a subsystem on the bus. Monitor 46 connects to the bus through Display Adapter 30. A relative pointing device (RPD) 48 such as a mouse connects through Serial Port 40. Some devices such as Keyboard 50 can communicate with the CPU by direct means without using the main data bus as, for example, via an interrupt controller and associated registers (not shown).

As with the external physical configuration shown in Fig. 2A, many subsystem configurations are possible. Fig. 2B is illustrative of but one suitable configuration. Subsystems, components or devices other than those shown in Fig. 2B can be added. A suitable computer system can be achieved without using all of the subsystems shown in Fig. 2B. For example, a standalone computer need not be coupled to a network so Network Interface 44 would not be required. Other subsystems such as a CDROM drive, graphics accelerator, etc. can be included in the configuration without affecting the performance of the system of the present invention.

Fig. 2C is a generalized diagram of a typical network.

In Fig. 2C, the network system 80 includes several local networks coupled to the Internet. Although specific network protocols, physical layers, topologies, and other network properties are presented herein, the present invention is suitable for use with any network.

In Fig. 2C, computer USER1 is connected to Server1. This connection can be by a network such as Ethernet, Asynchronous Transfer Mode, IEEE standard 1553 bus, modem connection, Universal Serial Bus, etc. The communication link need not be a wire but can be infrared, radio wave transmission, etc. Server1 is coupled to the Internet. The Internet is shown symbolically as a collection of server routers 82. Note that the use of the Internet for distribution or communication of information is not strictly necessary to practice the present invention but is merely used to illustrate a preferred embodiment, below. Further, the use of server computers and the designation of server and client machines is not crucial to an implementation of the present invention. USER1 Computer can be connected directly to the Internet. Server1's connection to the Internet is typically by a relatively high bandwidth transmission medium such as a T1 or T3 line.

Similarly, other computers at 84 are shown utilizing a local network at a different location from USER1 computer. The computers at 84 are coupled to the Internet via Server2. USER3 and Server3 represent yet a third installation.

Note that the concepts of "client" and "server," as used in this application and the industry, are very loosely defined and, in fact, are not fixed with respect to machines or software processes executing on the machines. Typically, a server is a machine or process that is providing information to another machine or process, i.e., the "client," that requests the information. In this respect, a computer or process can be acting as a client at one point in time (because it is requesting information) and can be acting as a server at another point in time (because it is providing information). Some computers are consistently referred to as "servers" because they usually act as a repository for a large amount of information that is often requested. For example, a World Wide Web (WWW, or simply, "Web") site is often hosted by a server computer with a large storage capacity, high-speed processor and Internet link having the ability to handle many high-bandwidth communication lines.

A server machine will most likely not be manually operated by a human user on a continual basis, but, instead, has software for constantly, and automatically, responding to information requests. On the other hand, some machines, such as desktop

computers, are typically thought of as client machines because they are primarily used to obtain information from the Internet for a user operating the machine.

Depending on the specific software executing at any point in time on these machines, the machine may actually be performing the role of a client or server, as the need may be. For example, a user's desktop computer can provide information to another desktop computer. Or a server may directly communicate with another server computer. Sometimes this is characterized as "peer-to-peer," communication. Although processes of the present invention, and the hardware executing the processes, may be characterized by language common to a discussion of the Internet (e.g., "client," "server," "peer") it should be apparent that software of the present invention can execute on any type of suitable hardware including networks other than the Internet.

Although software of the present invention, may be presented as a single entity, such software is readily able to be executed on multiple machines. That is, there may be multiple instances of a given software program, a single program may be executing on two or more processors in a distributed processing environment, parts of a single program may be executing on different physical machines, etc. Further, two different programs, such as a client and server program, can be executing in a single machine, or in different machines. A single program can be operating as a client for one information transaction and as a server for a different information transaction.

Fig. 3 shows a model chart for basic data objects of the present invention.

For example, Fig. 3 shows that a document object 200 includes a "BLOB_ID," "DESCRIPTION," "DOCUMENT_ID," etc. The BLOB_ID, in turn, references a blob table that includes additional information and/or references to other objects. Naturally, any manner of suitable software implementation can be employed that can use different, or varied, architectures than that of Fig. 3.

A preferred embodiment of the present invention operates on Microsoft NT or Unix platforms acting as servers. Client platforms can be any personal computer, consumer processing device, etc. The preferred embodiment uses Java for client-side processing and uses Oracle 8i for the underlying database engine. Again, given the nature of computer processing, these specifics are merely one implementation of hardware and software that can be used to implement the present invention.

Although the present invention has been discussed with respect to specific embodiments, these embodiments are merely illustrative, and not restrictive, of the invention. The scope of the invention is to be determined solely by the appended claims.

5

002250 94657560